



Demand Moderation in Military Communication Networks

Perry A. Blackmore

DSTO-TR-1286

DISTRIBUTION STATEMENT A

Approved for Public Release
Distribution Unlimited

20020729 207



Demand Moderation in Military Communication Networks

Perry A. Blackmore

Communications Division
Electronics and Surveillance Research Laboratory

DSTO-TR-1286

ABSTRACT

Interconnected heterogeneous networks with diverse ownership, carrying a wide range of multimedia traffic with extreme variations in load will characterise military networks of the future. Consequently, the design of a management and control architecture, which fosters efficient and military-tailored resource sharing, is a challenging multifaceted problem. A key feature of such an architecture is to give the network the ability to deliver meaningful signals to users in order for them to modify their behaviour in a way that is beneficial to the network as a whole. For example, during times of network stress, users should be discouraged from excessive network usage. Demand moderation is the term used to encompass the array of mechanisms aimed at achieving this end. Integrated Defence networks of the future should benefit enormously from demand moderation mechanisms.

RELEASE LIMITATION

Approved for public release

AQ F02-10-2488

Published by

*DSTO Electronics and Surveillance Research Laboratory
PO Box 1500
Edinburgh South Australia 5111 Australia*

*Telephone: (08) 8259 5555
Fax: (08) 8259 6567
© Commonwealth of Australia 2002
AR-012-176
April 2002*

APPROVED FOR PUBLIC RELEASE

Demand Moderation in Military Communication Networks

Executive Summary

In the next few years, one can envisage a highly connected and integrated Defence communications network comprising high speed land lines, satellite links, HF links, and a variety of other RF links. With this increased connectivity, there comes a wide array of obvious benefits. However, a potentially serious problem will be introduced, which must be addressed. Users who reside on high capacity networks can inadvertently overload low capacity or impoverished networks through inappropriate usage. There have been already been examples of this in recent times including the recent Interfet operations.

This report provides an introduction to demand moderation concepts that concern mechanisms aimed at controlling user behaviour on communications networks. The key idea is to provide feedback signals to users in order to modify behaviour. These signals can have the effect of encouraging user activity during times of low network load, and deterring it during overload. Intelligent modules placed in the communications networks are the means by which these signals are generated. These modules would be aware of traffic flows, local link conditions, and user demand. Advanced network modules could perform distributed scheduling in order to further co-ordinate network usage. Signals processed by user modules could provide users with concise state information via a graphical interface – hopefully moderating behaviour to benefit the enterprise.

This report describes the key elements of a demand moderation architecture. An example architecture is presented for the protection of a single impoverished link, which is a common scenario in Defence networks. An auction based token bucket algorithm is developed which is used to realise this architecture. This algorithm increases the cost of network usage according to the user demand and network load. It has been implemented in a demonstrator at DSTO Fernhill and is readily realisable in real-time.

The demand moderation concepts presented in this report, augment the early work done by Communications Division in the area of Military QoS (MQoS) within Defence networks. Additional mechanisms are described that enable the moderation of flows within a traffic class. These classes, described in the original framework, could possibly be vulnerable to inappropriate user access without moderation mechanisms.

Authors

Perry A. Blackmore Communications Division

Perry Blackmore completed a Bachelor of Science at the University of Melbourne in 1986, double majoring in Mathematics and Physics. He then completed a Graduate Diploma of Education in 1987 at the Melbourne College of Advanced Education. After 3 years of secondary teaching, he completed a Graduate Diploma of Science in Theoretical Physics at the Australian National University. His thesis that year was entitled "Design Parameter Selection in Linear Quadratic Gaussian Control". He then completed a PhD in control theory in 1995 with his thesis entitled "Discretization Methods for Control Systems Design". He is now employed as a Senior Research Scientist with Networks Architectures Group, Communications Division, DSTO. His recent research activities have included scheduling and protocol development for satellite communications, Quality of Service for military networks, and demand moderation for heterogeneous networks.

Contents

1. INTRODUCTION	1
2. USER INFLUENCE AND ARCHITECTURE REQUIREMENTS.....	3
3. AUCTION BASED TOKEN BUCKET ALGORITHM.....	6
3.1 Overview.....	6
3.2 Algorithm Description.....	8
3.2.1 Auction Thread	8
3.2.2 Real-Time Traffic Thread.....	12
4. EXAMPLE ARCHITECTURE.....	13
4.1 Description.....	13
4.2 Implementation Details.....	14
5. FURTHER WORK	15
6. CONCLUSION.....	18
7. REFERENCES.....	18

1. Introduction

In the last five years the number of Internet hosts has increased from approximately 5 million to 93 million, and the number of web sites from 20000 to 20 million [ZAKO00]. User expectations of the new technology have inflated with this explosive growth of information availability. Additionally, the furious development of new multimedia/multiparty services has further fuelled users' technological hunger.

The provisioning of a wide variety of services has revealed the inadequacies of the currently used Internet network protocol, IP Version 4, which only offers a best effort service class. Anyone who has attempted to operate a streamed audio application across the Internet will affirm this. Additionally, the best effort service offered by the current Internet does not encourage or enable sensible resource sharing. This has a parallel in history. For hundreds of years, herders grazed their cattle and sheep on common land. As long as no individual tried to graze too many cattle, everybody benefited from the common resource. The problem is that the benefit for taking more than their share went to the freeloading herder while the cost was paid by everyone. When too many people asserted their self-interest above the interest of the commons, overgrazing destroyed the value of the common land. Biologist Garrett Hardin described this catastrophe as "the tragedy of the commons."

The urgent need to support Quality of Service (QoS) has generated a great deal of effort from the Internet Engineering Task Force (IETF) who have proposed candidate QoS models such as the Integrated Services architecture [NWG94], and the Differentiated Services architecture [NWG98]. A well designed management and control architecture for integrated communications networks is paramount.

Paralleling the commercial situation, there is a growing demand amongst Defence users for multimedia/multiparty services on Defence networks. To meet this need, the Australian Defence Force (ADF) is assembling a fixed high bandwidth core network. Additionally, there will be improvements to the currently impoverished tactical communications infrastructure and its interconnections with the strategic domain. The creation of a highly connected homogeneous network infrastructure will result.

As argued in [KWIA99a, KWIA99b, BLAC00] there will be a strong need to have mechanisms in place to support Military Quality of Service (MQoS) on this infrastructure. MQoS extends QoS concepts by introducing the notion of *military value* of information. As argued in [KOWA96a, KOWA96b, KWIA99a, KWIA99b], network traffic should be managed as a function of its military value. In overloaded military networks when not enough resources are available to send all information at a desired QoS, flow prioritisation and graceful degradation should be based on the military value of the flow.

In [KWIA99b] a set of mechanisms are presented for achieving end-to-end MQoS. These are provision (flow establishment), control (real-time scale flow maintenance), and

management (slow time scale flow maintenance). The implementation of these will go a long way towards achieving the objective of realising interconnected networks which support a variety of services, and assure that the most important information based on military value is given preferential treatment.

Recent operational experiences have revealed shortcomings in our current networks that may not be completely resolved with MQoS mechanisms of priority and pre-emption. The following examples provide motivation.

- During the 1991 Gulf Crisis, military messages experienced enormous delays. In particular, FLASH messages were delivered well past their time of utility. Data gathered many years ago in [FELD79] indicate that under crisis conditions, as much as 60 percent of the traffic in certain military systems may be of precedence IMMEDIATE or higher. When such a high percentage of the traffic is urgent, priority queuing is not by itself adequate as a mechanism for managing congestion.
- During the Interfet operations in East Timor, a tactical extension to the SECBRS (Secret Backbone Router Service – secret strategic Defence network spread across Australia) was established between Australia and East Timor. The link provided a data channel at 256kbps. This link was quickly saturated by users accustomed to sending large email attachments across strategic networks. An interim procedure was eventually established which limited the size of items that could be transferred across the boundary.
- Again, during the Interfet operations, an experimental Theatre Broadcast System (TBS) was deployed. The system gave users the ability to request items to be included on the broadcast via a common interface. The interface enabled users to specify a number of delivery priority options. It was found that users who were dissatisfied with delivery times inflated their precedence and deflated their required delivery times. This behaviour was even seen by the system designers who should have known better!

Two major lessons can be drawn from the preceding examples. Firstly, during times of conflict, users have an increased belief in the importance of their information and justify their network usage accordingly. Secondly, with the expected increase in strategic network extensions to the tactical domain, users can, from the comfort of their office, access relatively impoverished communications links without necessarily appreciating the ramifications of their actions.

Clearly, mechanisms need to be in place in military networks to control user behaviour. These mechanisms should not only discourage or inhibit network usage in times of high load, but also encourage or facilitate network usage in times of low load. Further, we argue that intelligent network mechanisms can be used to modify user behaviour on military networks in a manner that results in a greater enterprise network utility. This approach shall be termed *Demand Moderation*.

At the outset, we deflect three arguments that can be raised against demand moderation concepts. The first is that increasing bandwidth will alleviate the problems identified. This

hardly deserves a response other than a reminder of economics, physics (Shannon's Law), and the Internet experience (bandwidth demand outpaces bandwidth supply - three years ago bandwidth demand doubled every year, now it doubles every 2-3 months).

A second argument is that an MQoS framework, such as the one advocated in [KWIA99a, BLAC00], is sufficient to deal with the identified issues. The MQoS framework suggests that the priority of an information flow in military communications networks should be managed as a function of its military value. This is based on: the *Mission Priority* value, which represents an enterprise value relative to other missions; the *Precedence* which is used to request from the network preference to transfer important information over a limited period of time; and the *User perceived priority* which allows the user to make a subjective differentiation between his/her applications and/or flows when all other parameters are equal. Such an architecture will indeed alleviate many of the problems identified by the examples above. However, in order to deal with the full array of problems, we argue that demand moderation should form a key component of this architecture. The Gulf Crisis and the TBS experience show that priority systems are open to user abuse. Users with high privilege can unknowingly cripple networks with inappropriate usage. Even inappropriate submissions of low priority/precedence traffic on networks can adversely affect low privilege users in an undesirable way. Further, groups of users limited to particular priority levels need to be moderated against one another.

Thirdly, the use of military procedures, policies and guidelines can enforce appropriate network behaviour. This in and of itself is a form of demand moderation. However, as traditionally enforced, such an approach is somewhat archaic on modern networks. Inefficiencies are created when simple policies (e.g., no more than 10 flash messages per day, or max length of a message is 60 characters) are enforced on human users. Rather, modern networks can have intelligence embedded in them with state information availability, which can trigger sophisticated algorithms that can deal with user requests in an automatic fashion. Consequently, a more efficient utilisation of the network by the enterprise may be achieved as opposed to a set of static policies.

The remainder of this paper is organised as follows. Section 2 will provide a survey of the requirements for a demand moderation architecture. An example architecture is presented in Section 3. A demand moderation demonstrator is described in Section 4. Finally, in Section 5, conclusions, remaining issues, and further work are given.

2. User Influence and Architecture Requirements

Demand moderation mechanisms attempt to exert a control upon and modify user network behaviour. This is to be distinguished from congestion and control mechanisms that directly shape network traffic. However, the latter mechanisms may form part of a demand moderation architecture. Network bandwidth, a scarce resource on many military communications networks is to be used wisely by the enterprise as a whole. "Rogue" users

are to be discouraged and network usage aligned with enterprise objectives should be encouraged.

The question of allocating and controlling a scarce resource is of course a fundamental economics problem for which a wide range of economic systems have been devised. Whatever the economic system, four means have been identified by which decision-makers can motivate others to follow their pronouncements:

1. appeal to egocentrism through accumulation of wealth or respect of peers ("The user who can perform their required work while generated the least network traffic will be awarded the divisional prize!"),
2. reliance on tradition or customary obligation, ("In this organisation, we pride ourselves on generating low levels of network traffic!")
3. appeal to solidarity or the willing subordination of the individual's personal objectives to those shared by the group ("Think of your colleagues before you begin that download!"), and
4. coercion ("Anyone found generating high levels of network traffic will have the access removed!").

There is a range of ways to perform economic system taxonomies but arguably the most important is the level of centralisation. Centralisation refers to the extent allocation decisions are administered by central authorities or are taken by the micro-units independently of central authorities through a system of markets. Pure socialism and pure capitalism fall at the extremes of the centralisation spectrum.

In order for a centralised economy to be effective, the central decision making authority requires detailed and accurate state information. In [SEMR99a], where the focus is on pricing mechanisms in commercial networks to promote appropriate network sharing, the observation is made that centralised network control is not possible. This pronouncement is made on the basis of two observations. Firstly, because of the diverse ownership of network resources. Secondly, and more fundamentally, a large multimedia network with QoS would be too complex and decentralised to be under a single authority. Even centralised control objectives would be impossible to define. It is well accepted that network-wide optimisation is an impractical paradigm for the control of modern networks [ZHAN92, ALTM94, ORDA93, KORI95a, KORI95b]. Indeed, control decisions in large-scale networks are often made by each user independently, according to his own individual objectives. Such networks are defined as non-cooperative, and game theory is the appropriate framework to study their behaviour. Under the non-cooperative paradigm, the network is considered as a site of competition for the network resources among selfish users. The problem of selfish users on a network is not necessarily due to a pessimistic view, but rather recognising that technology and nature make it impossible to be unselfish, or even know what it is to be unselfish.

It is likely that decentralised or distributed market mechanisms are going to be more effective in controlling user behaviour, given the nature resource sharing in multimedia networks. Adam Smith's analysis of markets in the late 18th century showed that the

"invisible hand" effect of "selfish" distributed self-optimising agents often perform resource allocation more efficiently than centralised systems. Again in [SEMR99a] it is shown that by placing "rules" of interaction in the way the network is utilised, network wide optimisation can be performed. By taking a game theoretic approach in the engineering of the network, one can design mechanisms where the intelligent decision-making is distributed and the objective of a more efficient and fair utilisation of shared resources results from the induced dynamics. Using this approach, prices play a fundamental role as resource allocation signals.

From this, one can identify a number of key elements that must be present in a network in order for distributed market mechanisms to be effective in user control.

Users: The objective of demand moderation mechanisms is to exert a controlling influence on the individual network users.

Units: These are groups of one or more users who share common resources. The demand moderating mechanisms in the network treat units as atomic entities. Individual users within a group would be expected to be colocated and be able to negotiate unit resources.

Unit Funds: Each unit maintains a medium of exchange that can be traded for network usage.

Unit Income: There must be a means by which unit funds are generated.

Unit Agent: Each unit must have a representative that is responsible for interfacing with the network to trade unit funds for actual network resource. There must be a means by which units can instruct their agents to act according to their wishes. Once these instructions are issued, the agent, as far as possible, should act independently from the units.

Market: Competing unit agents require a forum for trading. The market is a virtual entity that brings unit agents together to negotiate for network resources.

Pricing: Prices are the terms of trade for exchanging unit funds for network resources in markets. Depending upon the exact demand moderation architecture, prices could relate to generic network usage, link-by-link usage, or bottleneck usage only.

Price Dynamics: In order to achieve the objectives of demand moderation under the economic model proposed, it is important that prices are driven by: unit demand for network resources; and network load.

Price Regulation: Prices are driven entirely by market forces in an idealistic market economy. In practical settings, there are other price regulating entities (e.g. government imposed price ceilings). In order to drive network dynamics to desirable states, mechanisms are required to monitor and regulate prices.

Network Monitoring: At key points in the network, especially at bottlenecks, state should be measured and be signalled via the price regulator to the units.

Signalling: There must be a means to signal prices to the various entities and players in the network.

Unit Interface: Users must be aware of the network dynamics, prices, and their own resources. The interface should be simple but effective in achieving this end.

Policing: Ultimately, there must be mechanisms to enforce behaviour and strong feedback to non-compliant units.

In the military context, demand moderation is crucial to cope with bottlenecks in impoverished links or overloaded networks. Demand moderation architectures designed to diminish bottlenecks would impose a high price for bottleneck traffic during times of stress. The price may be expected to rise purely through market forces, but also price-regulating mechanisms could enforce this as mentioned above. The fundamental philosophy of demand moderation is that these prices must be available to individual units in order to modify their network usage accordingly. Network usage during high price periods would result in rapid loss of funds and finally may result in access denial through either an administrative process, or an automated regime. Conversely, during times of lower prices, users may be encouraged to use the network.

Any demand moderation architecture requires some careful decisions in relation to scalability, particularly when it is network-wide. The signalling overhead must be minimized by keeping message sizes small and by minimizing inter-network communications, particularly those over bottlenecks that are trying to be protected. Trading off against this, computational efforts should be distributed where possible.

A further dimension that could be added to a demand moderation architecture is that of unit scheduling. In order to encourage considered network usage, it may be possible to announce ahead of time particular periods of economical network usage. This could take the form of price reductions over a period of time, or increased income for the unit.

This report will not attempt to specify a generic demand moderation architecture. However, in the following sections, a simple demand moderation architecture is proposed for the protection of a single impoverished link. In order to realise this architecture, the next section introduces the auction based token bucket algorithm. Further work needs to be performed to determine the efficacy of particular schemes.

3. Auction Based Token Bucket Algorithm

3.1 Overview

This section presents an auction based token bucket algorithm that is designed to meet the demand moderation requirements for a single communications link. The role of the algorithm is not limited to demand moderation - there may be tolling applications in the commercial setting. It should be stressed that no claims are made for the uniqueness of the algorithm - there may be other means to achieving a demand moderation utility. The key features are:

1. Bandwidth available for units is competed for in a market based on the Progressive Second Price (PSP) auction to be detailed below.

2. The key requirements for demand moderation, namely a concept of unit funds and of market prices, fit naturally into the algorithm framework.
3. The algorithm is split into two threads to render it realisable in real-time. It does not introduce significant latency to traffic streams.
4. Units can readily specify their auctioning "aggressiveness", sacrificing available funds for short-term network access if necessary.

The auction based token bucket algorithm is a variation of an algorithm in [KENN94] which was given a generalised treatment in [FLOY95]. The work of the latter reference forms a basis for the Class Based Queueing new networking features in the Linux 2.2 kernels.




Figure 1 - Auction Based Token Bucket Algorithm

A diagrammatic representation of the algorithm is given in Figure 1. Each unit has three associated resources:

1. A source of funds managed via a funds token bucket. Funds are deposited at a constant rate according to the token bucket attributes.
2. A bidder object which draws on available funds to participate in the auction. The bidders compete in order to obtain transmission tokens for their corresponding unit.
3. A transmission token bucket which is used to regulate unit traffic. Tokens are placed in the various buckets according to the auction results. Submitted unit traffic draws on the tokens for transmission.

An auctioneer object is responsible for maintaining transmission tokens, conducting auctions for these tokens, and for distributing those tokens to the successful bidder(s). During a particular auction, the various active bidders submit a series of (token quantity, bid price) bids based upon a pre-configured strategy/algorithm. These bids continue until either the auction reaches equilibrium (this is discussed in the next section) or until the auctioneer announces the completion to the various bidders. At the point of auction completion, transmission tokens are allocated according to the PSP auction allocation rule [LAZA98]. These tokens can then be drawn on by unit traffic.

The algorithm contains two independent threads: the auction thread and the real-time traffic thread. The auction thread manages the bidding, the PSP algorithm, and the distribution of the transmission tokens to the various transmission token buckets. The real-time traffic thread handles the packet scheduling and the release of transmission tokens required to send a packet.

The split nature of the algorithm is critical for the implementation of scheme as the introduction of excessive latency is minimised. The auction thread is generally computationally demanding, and if the dynamic auctioning scheme suggested by [LAZA98] is used, the convergence time is typically slower than traffic time scales. Decoupling the traffic thread allows packet scheduling independent of the state of the auction. It is not overly deleterious if the auctioning algorithm has not converged by the allocation and token distribution time. The allocation based on non-equilibrium bids has been found through experience to closely reflect the converged state allocation. This is due to the fact that bids generally converge rapidly to the neighbourhood of the Nash equilibrium (for a description of game theory and the concept of the Nash equilibrium the interested reader is referred to [FUDE91]), and then oscillate around it before ultimate convergence.

3.2 Algorithm Description

3.2.1 Auction Thread

Let auction allocations be performed at times $t_1, t_2, \dots, t_k, \dots$. Let the state of the auctioneer token bucket immediately subsequent to token allocation at time t_k be represented by $[\alpha_k^a, \beta^a, \rho^a]$, where α_k^a is the quantity of tokens in the bucket (which can take on a negative value), β^a is the capacity of the bucket and ρ^a is the rate at which tokens are placed in the bucket. ρ^a is generally set equal to the output line rate.

Let each funds token bucket be characterised by $[\alpha_{i,k}^f, \beta_i^f, \rho_i^f]$ for $i = 1, \dots, N$, where $\alpha_{i,k}^f$ is the quantity of tokens in bucket i immediately subsequent to t_k , β_i^f is the capacity of the funds bucket i , ρ_i^f is the constant rate at which tokens are placed in bucket i , and N is the number of bidders.

Let each transmission token bucket be characterised by $[\alpha'_{i,k}, \beta'_i, \rho'_{i,k}]$ where $\alpha'_{i,k}$ is the quantity of tokens in bucket i immediately subsequent to t_k , β'_i is the capacity of the output bucket i , and $\rho'_{i,k}$ is the number of transmission tokens (measured in bits) allocated to bucket i at time t_k .

Summed over all auctions, the total amount of tokens allocated to the transmission token buckets is less than or equal to the amount auctioned, i.e.

$$\sum_{k=1}^{\infty} \sum_{i=1}^N \rho'_{i,k} \leq \sum_{k=1}^{\infty} (t_k - t_{k-1}) \rho^a \quad (1)$$

Let n_k be the number of tokens allocated at an auction at time t_k . For implementation simplicity, assume n_k is constant $\forall k$. The average time between auctions (repletion time) is given by $\Delta t = n_k / \rho^a$.

Let t_c denote the current time and let t_l denote the last time when the auctioneer token bucket credit was updated. The auctioning algorithm is now presented.

Auctioning Algorithm

1. Set $k = 1$; $t_k = 0$; $t_l = 0$; $\alpha_k^a = 0$; $\alpha_{i,k}^f = \beta_i^f$ for $i = 1, \dots, N$
2. Update t_c
3. Update auctioneer token bucket credit: $\alpha_k^a = \min\{\alpha_k^a + (t_c - t_l)\rho^a, \beta^a\}$
4. If $\alpha_k^a \geq 0$ then allocate n_k transmission tokens (via equations (2) - (8) following);
 $\alpha_{k+1}^a = \alpha_k^a - n_k$; $k = k + 1$
5. Update fund buckets: $\alpha_{i,k}^f = \min\{\alpha_{i,k}^f + (t_c - t_l)\rho_i^f, \beta_i^f\} \quad \forall i$
6. $t_l = t_c$
7. Ask for bids; let n denote the number of new bids
8. If $n = 0$ then wait $\max\{-\alpha_k^a / \rho^a, 0\}$
9. Go to step 2

Progressive Second Price Auction

Step 3 of the auctioning algorithm involves the allocation of tokens to the various transmission token buckets associated with the bidders. The allocation is based on the current bid prices and bid quantities. The algorithm used is the Progressive Second Price (PSP) auction [LAZA00], which is an extension of the Vickrey or second price auction. The second price auction uses a closed ballot for auctioning a single indivisible object. The winner of the auction is the bidder with the highest bid, but only the second highest bid

price is paid. This method has many desirable properties, the most important of which is that it has an equilibrium profile where all participants bid their true valuation. The PSP auction, on the other hand, is concerned with the auctioning of a divisible object. The highest bidder receives his allocation first with a cost determined by a weighted sum of the lower bids. Any leftovers are then allocated to the second highest bidder who pays a weighted sum of the lower bids etc.

The description of the PSP auction to follow is a summary and tailoring of that presented in [LAZA00].

Let n_k tokens be allocated at time t_k . Let the bid of the i^{th} bidder just before the auction at time t_k be given by

$$s_{i,k} = (q_{i,k}, p_{i,k}) \quad (2)$$

where $q_{i,k}$ is the desired number of tokens and $p_{i,k}$ is the unit bid price. The set of all bid prices is denoted by

$$s_k = (s_{1,k}, \dots, s_{N,k}) = ((q_{1,k}, p_{1,k}), \dots, (q_{N,k}, p_{N,k})) \quad (3)$$

The set of all bids just before time t_k , apart from the i^{th} player is denoted by $s_{-i,k}$, i.e.

$$s_{-i,k} = (s_{1,k}, \dots, s_{i-1,k}, s_{i+1,k}, \dots, s_{N,k}) \quad (4)$$

Consequently one can denote the set of all bids by

$$s_k = (s_{i,k}; s_{-i,k}) \quad (5)$$

which emphasises that the i^{th} player is the focus.

The amount of tokens allocated to each bidder is given by

$$a_i(s_k) = \min\{q_{i,k}, \underline{Q}_i(p_{i,k}, s_{-i,k})\} \quad (6)$$

where

$$\underline{Q}_i(y, s_{-i,k}) = \max\left\{n_k - \sum_{p_j \geq y, j \neq i} q_{j,k}, 0\right\} \quad (7)$$

The cost of tokens allocated to the bidder is given by

$$c_i(s_k) = \sum_{j \neq i} p_{j,k} [a_j(0; s_{-i,k}) - a_j(s_{i,k}; s_{-i,k})] \quad (8)$$

Bidder Algorithm

In step 6 of the auctioning algorithm, each bidder is asked to update their bids. There are a wide range of possible methods that could be used and we now present a tailoring of an algorithm presented in [SEMR99a].

Each bidder i maintains a valuation function $\theta_i(z)$, relating allocation to benefit. The function should have a number of properties to reflect elastic demand, that is: $\theta_i(z)$

differentiable, $\theta'_i(z)$ positive, non-increasing, continuous, and strictly concave over a domain. These properties are needed to prove the convergence of the bidding algorithm to a Nash equilibrium.

Each bidder aims to maximise its utility, which is the difference between the valuation of the token allocation and the resulting cost of those tokens. To achieve this end, a utility function $u_i()$ is defined. Given the current opponent bids, i.e. $s_{-i,k}$, $u_i(s_k)$ can be calculated from the following

$$u_i(s_k) = \theta_i(a_i(s_k)) - c_i(s_k) \quad (9)$$

Each bidder maintains a "budget" denoted by $b_{i,k}$. This budget is the available funds to use in the bidding process. Naturally we have $b_{i,k} \leq \alpha_{i,k}^f \leq \beta_i^f$. The bidder must ensure that the resulting allocation cost of their bid does not exceed $b_{i,k}$.

The bid which maximises the utility function (9) at time t_k is given by

$$m_{i,k} = (v_{i,k}, w_{i,k}) \quad (10)$$

where

$$v_{i,k} = \left[\sup G_i(s_{-i,k}) - \varepsilon / \theta'_i(0) \right]^+ \quad \text{and} \quad w_{i,k} = \theta'_i(v_{i,k}) \quad (11)$$

with

$$G_i(s_{-i,k}) = \left\{ z \in [0, n_k] : z \leq Q_i(\theta'_i(z), s_{-i,k}) \quad \text{and} \quad \int_0^z P_i(\eta, s_{-i,k}) d\eta \leq b_{i,k} \right\} \quad (12)$$

and

$$P_i(z, s_{-i,k}) = \inf \{ y \geq 0 : Q_i(y, s_{-i,k}) \geq z \} \quad (13)$$

and

$$Q_i(y, s_{-i,k}) = \lim_{\eta \downarrow y} Q_i(\eta, s_{-i,k}) = \max \left\{ n_k - \sum_{p_{j,k} > y, j \neq i} q_{j,k}, 0 \right\} \quad (14)$$

A cut-off parameter ε is introduced to expedite the convergence of the bidding algorithm. A bidder will only update the current bid if the utility is increased by at least ε .

Here is the bidder algorithm.

1. If $\alpha_{i,k}^o = \beta_i^o$ (output token bucket full) return no bid
2. Set auction budget $b_{i,k} = \frac{1}{2} \alpha_{i,k}^f$
3. Compute best bid $m_{i,k}$ according to (10)
4. If $u_i(m_{i,k}; s_{-i,k}) > u_i(s_{i,k}; s_{-i,k}) + \varepsilon$ then return $m_{i,k}$
5. Return no bid

In step 2, the budget is always set to half the available funds. If one assumes a constant inter-auction time of $\Delta t = t_{k+1} - t_k$, then one can show as the number of auctions n grows, that the auctioning budget relating to the i^{th} bidder is given by

$$\lim_{n \rightarrow \infty} \left\{ \frac{1}{2^n} \alpha_{i,0}^f + \Delta t \rho_i^f \sum_{i=0}^n \frac{1}{2^n} \right\} = 2 \Delta t \rho_i^f \quad (15)$$

In other words, assuming steady state, one can expect that the amount of funds used in each auction is given by $\Delta t \rho_i^f$ (c.f. step 2 of the bidding algorithm).

An example valuation function that can be used is given by

$$\theta_i(z) = -\kappa_i (\min\{z, \bar{n}_i\})^2 / 2 + \kappa_i \bar{n}_i (\min\{z, \bar{n}_i\}) \quad (16)$$

with $0 < \bar{n}_i < n_k$. The parameters κ_i and \bar{n}_i can be configured by the unit to reflect the bidding strategy.

3.2.2 Real-Time Traffic Thread

Unlike many other token bucket algorithms that have a constant fill rate, the auction based token bucket algorithm is serviced at a non-constant rate that depends on auction results. The traffic service algorithm is fairly simple. Let $l_{i,k}$ denote the length of packet k , $k = 1, 2, \dots$, in buffer i , $i = 1, 2, \dots, N$. $l_{i,1}$ represents the length of the first packet in buffer i . $l_{i,1} = 0$ if and only if buffer i is empty. Let π_i denote a buffer metric used for scheduling prioritisation. The packet scheduling algorithm operates as follows.

Packet Scheduling Algorithm

1. For each i set $\pi_i = -\infty$
2. For each i , if $l_{i,1} > 0$ then $\pi_i = \alpha_{i,k}^f - l_{i,1}$ (the difference between the transmission token bucket credit and the length of the first packet)
3. Let $\bar{\pi} = \max \pi_i$
4. If $\bar{\pi} > 0$ then send first packet from buffer i ($i : \arg \pi_i = \bar{\pi}$)
5. Go to step 1

The above algorithm offers a guaranteed service to any streams with sufficient credit. in their transmission token bucket. Conversely, packets will be delayed if their bidder has not accrued sufficient credit via the auctions.

The algorithm can be modified so that tokens that are not allocated in any auction (recall equation (1)) are distributed (free of charge) via some form of round robin scheme. This would allow units with insufficient funds to capture any leftover or unguaranteed bandwidth.

4. Example Architecture

In this section, an example demand moderation architecture will be presented. The architecture is a limited one which is limited to moderation of user traffic over a single impoverished link. Despite the simplicity, this situation is ubiquitous in Defence networks. Some examples include:

Local Area Network Connected to the Wide Area Network via a Narrowband Link: This situation is quite common. Many Defence bases are in this situation.

Deployed Forces Local Area Network: Brigade or Battalion level forces are typically connect to headquarters via satellite links (e.g. Parakeet). These links are limited to data rates in to order of 64-512kbps.

VSAT Networks: In the near future, the Australian Defence Force will be acquiring a payload on a commercial satellite. One of the possible uses for this payload is to support a multiple access network with multiple distributed users competing for common bandwidth.

Remote Servers: Servers on termination of impoverished links fit into this model.

Bottlenecks: It may be possible to abstract certain networks into a single bottleneck (Norton equivalent bandwidth [HSIAO89]) capacity and moderate the users based on that abstraction.

4.1 Description

Figure 2 represents a group of users sharing a common local area network that is connected to a wide area network via a single impoverished link. In the example diagram, there are three units with the first unit containing multiple users sharing common demand moderating resources.

In this architecture, the auction based token bucket algorithm presented in the previous section is used as the demand moderating market mechanism. Each unit has an agent in the form of a bidder, which may reside in either a unit workstation or downloaded onto the gateway where the auctioneer resides. Bids are submitted asynchronously either across the LAN or via inter-process communications. An auctioneer receives bids from the bidders and transmission tokens are allocated to the various units, which enables access to link. While the state of the auction is a complex one, only the current unit funds and the high bid price are fed back to the users.

Generally, the entire bandwidth is to be allocated as a result of the auction (this may be shared among the units). However, it is possible to limit the traffic entering the WAN by feeding back a signal to the auctioneer specifying the auction reserve price. This may have the effect of only a partial allocation of the available outgoing bandwidth. This mechanism can be used to control the wide area dynamics. That is why a

dummy bidder is normally configured into the auctioneer. This bidder submits bids at all auctions asking for all the tokens at some "reserve price". In order to obtain any tokens via the Progressive Second Price algorithm, the other bidders must bid at a price higher than this reserve price.




Figure 2- Demand Moderation for a Single Impoverished Link

As the unit's funds are diminished, the bidder will be unable to effectively compete in the auction with the result that the unit's QoS will be compromised.

4.2 Implementation Details

A demonstrator has been developed to explore the efficacy of the preceding architecture in terms of user influence. A demand moderating router, which runs the auctions, has been developed which operates on a workstation running the Linux 2.4.2 operating system enabled with *netfilter*¹. The router sits on a private local area network which connects to a wide area network via an emulated narrow-band link.

¹ The *netfilter* framework was introduced in the Linux 2.4 kernels and was designed to simplify datagram processing in the kernel firewalling code. The resulting architecture is far more flexible compared to previous versions. *Netfilter* provides a mechanism for passing packets out of the stack for queueing to userspace, then receiving these packets back into the kernel with a verdict specifying what to do with the packets

Iptables is used to set up, maintain, and inspect the tables of IP packet filter rules in the Linux kernel. For our purposes, the Linux workstation is configured by iptables to operate as a masquerading router, and additionally, to pass network packets into the userspace. Using **libipq**², packets that are redirected into userspace are queued and serviced according to auction outcomes before being reinserted into the kernel.

The following commands are executed as root to enable the routing and userspace packet delivery.

```
# Turn on masquerading
echo "1" > /proc/sys/net/ipv4/ip_forward
/sbin/iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
# Enable packet forwarding.
iptables -A FORWARD -j QUEUE
```

Experiments remain to be performed using real users to determine the efficacy of this architecture and approach.

5. Further Work

Demand moderation forms part of the wider network management and control effort by Communications Division, DSTO. Work in the division has focussed on architectures supporting Military Quality of Service (M-QoS). Initial work had a long-term focus and was directed towards a network control and management framework based on models devised by the IEEE standard initiative called P1520 and the TINA-C consortium. This framework assumes the use of a network environment that is open and programmable. The framework assumes that each switch/router or end-station device advertises its resources, making them accessible to controllers which run algorithms on a general purpose distributed processing environment.

Recent work has focussed on two areas: the M-QoS interface, and M-QoS in IP networks. To provide M-QoS, a standard interface, called the M-QoS interface, between user applications and network management and control has been developed. This facilitates application specification of M-QoS requirements for a network, as well as enabling the network to inform the application about any problems in delivering promised level of service. The interface assumes two sets of parameters, one describing commercial QoS parameters and the other, military specific parameters. These are used to evaluate the ultimate priority of the flow according to the policy and associated algorithm currently implemented in the network.

Figure 3 demonstrates the approach to achieve the M-QoS capability in a Defence network using network management and control mechanisms. The network control and management evaluates the priority of the flow based on the provided parameters and the

². The development library libipq is a development library for iptables userspace packet queuing.

currently implemented enterprise policy. If there are enough resources in the network to establish the flow for this priority level, the user is notified and the network management establishes the flow with the corresponding priority.




Figure 3 - MQoS Representation

Two recent reports ([KWIA01a], [KWIA01b]) have focussed on providing MQoS in IP networks. Figure 4 presents the current view on how M-QoS could be implemented in Defence networks, both strategic and tactical. The use of DiffServ, Bandwidth Brokerage, MPLS and IETF policy framework is assumed. The whole Defence communication infrastructure is divided into strategic and tactical domains (more than one of each type may coexist). Policy Administration (PA) is responsible for consistent DiffServ offering across all Defence Core domains. It controls the various Bandwidth Broker (BB) and Policy Decision Points (PDP) by distributing policy changes and receiving feedback about the state of the network. As highlighted in [KWIA01b], there are arguments for distributing the PA function that will not be discussed here. Each domain is equipped with a BB/PDP which interact with the PA, and perform bandwidth broker functions such as: admission control of requests coming from end-user hosts for network resources; evaluation of the ultimate priority of the flow using an algorithm distributed by the PA; and selection of a particular DiffServ class that should accommodate a flow (if there is enough resources within the class, the request is admitted and the edge routers are notified how to mark the flow). Note that BBs of adjacent domains need to cooperate in order to satisfy cross domain flow requests.

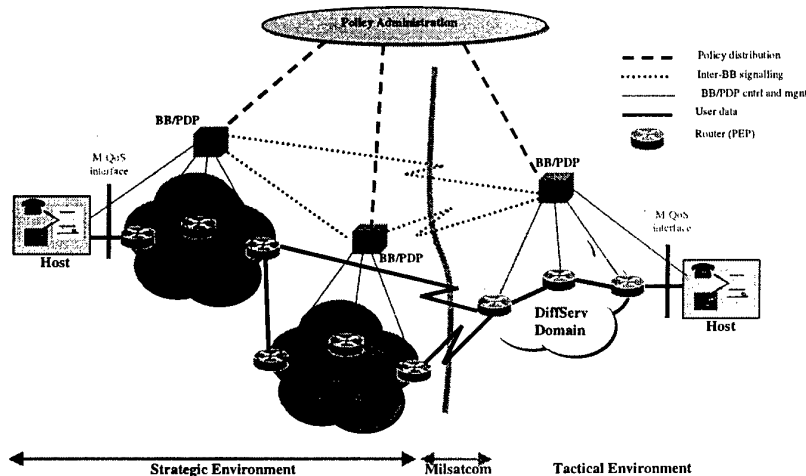


Figure 4 – Proposed MQoS Framework (BB – Bandwidth Broker, PDP – Policy Decision Point, PEP – Policy Enforcement Point)

Demand moderation mechanisms can be synthesised into this framework within classes of flows representing particular priority levels. Moderation is particularly needed in the tactical domains and part of the brokerage function could act as a source of price regulations to achieve the desired wide area dynamics. Further, in the flow admission process, unit funds could be taken into account. This principle could also be used in the degradation of existing streams.

A major Communications Division MILSATCOM task, Multiple User Access Architectures, is currently in the preparation phase. This will provide an ideal environment to explore the integration of MQoS and demand moderation concepts. A multiple user architecture with a central controller could readily incorporate the sort of demand moderation scheme presented in the previous section, with the auctioneer co-located with the network controller. The BB/PDP could also be integrated with the network controller.

There are a number of significant challenges remaining before demand moderation concepts could be widely integrated into Defence networks. Demand moderation is focussed upon influencing users as they interact with the network. Generally this relies on user traffic being able to be measured. However a significant proportion of user traffic is likely to be generated via proxies and servers. There is the issue of “charging” for both “pushed” and “pulled” data. Furthermore, encryption means that only edge traffic can be observed. Compounding this, on many networks most of the traffic is generated independently from human users and demand moderation may have little overall benefit.

This paper has presented an architecture based on a single impoverished link. There remains the challenge of moderation across generic networks. There is the appealing prospect of shaping wide area dynamics via the distributed market mechanisms as hinted in [SEMR99a].

Of course the wider application of demand moderation will most likely require the fitting of edge routers with custom algorithms. This may be prohibitive and adjunct processors may be fitting only where moderation is particularly critical.

6. Conclusion

Demand moderation is an area of work that is likely to reap significant benefits for many Defence networks. With heightening user expectations of communications networks and the simultaneous expansion of Defence networks into the tactical arena, demand moderation mechanisms will form an important component of future networks.

The author of this report has performed an initial survey of the area and has presented a simple architecture for the moderation of a single impoverished link. An auction based token bucket algorithm has been presented to realise this architecture.

The demand moderation concepts presented here, augment the early work done by the Communications Division of DSTO in the area of Military QoS (MQoS) within Defence networks. Additional mechanisms have been described in this report that enable the moderation of flows within a traffic class. These classes, described in the original framework, could possibly be vulnerable to inappropriate user access without moderation mechanisms.

There remain a number of significant challenges and much further work in this area. A generic architecture needs to be developed which is scalable and can be realised without large signalling overheads. The contexts most suitable for demand moderation mechanisms need to be identified. Further work will be performed to further integrate the moderation concepts into the MQoS framework. Finally user trials are required to quantify the likely effects of moderating schemes.

7. References

- [ALTM94] E. Altman, "Flow Control Using the Theory of Zero-Sum Markov Games", IEEE Transactions on Automatic Control, 39(4), pp. 814-818, April 1994.
- [BLAC00a] P. Blackmore, P. George, M. Kwiatkowski "A Quality of Service Interface for Military Applications", proceedings MILCOM 2000 conference, Los Angeles, Oct 2000.

- [FELD79] N. Feldman, W. Sollfrey, S. Katz, and S. Dudzinsky Jr., "Writer-to-Reader Delays in Military Communications Systems", RAND, R-2473-AF, Santa Monica, California, October 1979.
- [FLOY95] S. Floyd and V. Jacobson, "Link-sharing and Resource Management Models for Packet Networks", IEEE/ACM Transactions on Networking, Vol.3, No.4, 1995.
- [FUDE91] D. Fudenberg and J. Tirole, "Game Theory", MIT Press, 1991.
- [HSIAO89] M. Hsiao and A. Lazar, "An Extension to Norton's Equivalent", Queueing Systems, Theory and Applications, 5:401-412, 1989.
- [KENN94] A. U. Kennington, "Credit priority queueing for ATM switches", ATNAC 1994 conference, Melbourne, December 1994, pp. 399-403.
- [KORI95a] Y. Korilis and A. Lazar, "On the Existence of Equilibria in Noncooperative Optimal Flow Control", Journal of the ACM, 42(3), pp. 584-613, May 1995.
- [KORI95b] Y. Korilis, A. Lazar, and A. Orda, "Architecting Noncooperative Networks", IEEE Journal on Selected Areas in Communications, 13(7), pp. 1241-1251, September 1995.
- [KOWA96a] N. Kowalchuk, P. Blackmore, "Multimedia Services in a Tactical Environment", MILCOM'96, USA, 1996.
- [KOWA96b] N. Kowalchuk, J. Sholtz, T. Moors, "A Framework for the Management of Networks Based on Information Value", MILCOM'96, USA, 1996.
- [KWIA99a] M. Kwiatkowski, P. George, "A Network Control and Management Framework Supporting Military Quality of Service", MILCOM'99, Atlantic City, USA, 1999.
- [KWIA99b] M. Kwiatkowski, "Network Control and Management Architectural Framework Supporting Military Quality of Service", DSTO-TR-0871, 1999.
- [KWIA01a] M. Kwiatkowski, "Preliminary Assessment of Transmission Technologies to Support Military Oriented QoS", DSTO-TR-1207, 2001.
- [KWIA01b] M. Kwiatkowski, "A Concept of the Defence Core Communication Infrastructure Supporting M-QoS", DSTO Technical Report DSTO-TR-1220, January 2002.
- [LAZA00] A. Lazar, N. Semret, "Design and Analysis of the Progressive Second Price Auction for Network Bandwidth Sharing", Telecommunications Systems, Special Issue on Network Economics, to appear.
- [LAZA98] A. Lazar, N. Semret, "The Progressive Second Price Auction Mechanism for Network Resource Sharing", 8th International Symposium on Dynamic Games, Maastricht, The Netherlands, July 5-8, 1998.
- [NWG94] R. Braden, D. Clark, S. Shenker, "Integrated Services in the Internet Architecture: an Overview", Internet Draft, June 1994.
- [NWG98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, "RFC 2475 - An Architecture for Differentiated Services", Internet Draft, December 1998.

- [ORDA93] A. Orda, R. Rom, and N. Shimkin, "Competitive Routing in Multiuser Communication Networks", *IEEE/ACM Transactions on Networking*, 1(5), pp. 510-521, October 1993.
- [SEMR00] N. Semret, R. Liao, A. Campbell, A. Lazar, "Peering and Provisioning of Differentiated Internet Services", *IEEE Infocom 2000*, Tel Aviv, Israel, March 2000.
- [SEMR99a] N. Semret, "Market Mechanisms for Network Resource Sharing", PhD Thesis, Columbia University, 1999.
- [SEMR99b] N. Semret, A. Lazar, "Spot and Derivative Markets in Admission Control", *16th International Teletraffic Congress*, Edinburgh, UK, June 1999.
- [SEMR00] N. Semret, R. Liao, A. Campbell, A. Lazar, "Market Pricing of Differentiated Internet Services", *7th IWQoS*, London, May 31 - June 4, 1999.
- [ZAKO00] Robert H'obbes' Zakon, "Hobbes' Internet Timeline", <http://www.isoc.org/guest/zakon/Internet/History/HIT.html>, Copyright (c) 1993-2000 by Robert H Zakon.
- [ZHAN92] Z. Zhang and C. Douligeris, "Convergence of Synchronous and Asynchronous Greedy Algorithms in a Multiclass Telecommunications Environment", *IEEE Transactions on Computers*, 40(8): pp. 1277-1281, August 1992.

DISTRIBUTION LIST

Demand Moderation in Military Communication Networks

Perry A. Blackmore

AUSTRALIA

DEFENCE ORGANISATION

Task Sponsor

Director General Command, Control, Communications and Computers (DGC4)
PD JP 2047 (CWAN)
PD JP 2068 (NOC)
Head Defence Information Systems Group

S&T Program

S&T Program

Chief Defence Scientist	} shared copy
FAS Science Policy	
AS Science Corporate Management	
Director General Science Policy Development	
Counsellor Defence Science, London (Doc Data Sheet)	
Counsellor Defence Science, Washington (Doc Data Sheet)	
Scientific Adviser to MRDC Thailand (Doc Data Sheet)	
Scientific Adviser Joint	
Navy Scientific Adviser (Doc Data Sheet and distribution list only)	
Scientific Adviser - Army (Doc Data Sheet and distribution list only)	
Air Force Scientific Adviser	
Director Trials	

Aeronautical and Maritime Research Laboratory

Director

Electronics and Surveillance Research Laboratory

Chief of Communications Division
Research Leader Military Information Networks (Dr. M. Flaherty)
Head Network Architectures (Dr. J. Asenstorfer)
Task Manager: Dr. M. Kwiatkowski
Author: Dr. P. Blackmore (2 copies)
Dr. P. Shoubridge
Mr. P. George

DSTO Library and Archives

Library Salisbury (2 copies)
Australian Archives

US Defense Technical Information Center, 2 copies
UK Defence Research Information Centre, 2 copies

Canada Defence Scientific Information Service, 1 copy
NZ Defence Information Centre, 1 copy
National Library of Australia, 1 copy

Capability Systems Staff

Director General Maritime Development (Doc Data Sheet only)
Director General Aerospace Development (Doc Data Sheet only)

Army

Stuart Schnaars, ABCA Standardisation Officer, Tobruck Barracks, Puckapunyal,
3662(4 copies)
SO (Science), Deployable Joint Force Headquarters (DJFHQ) (L), MILPO Gallipoli
Barracks, Enoggera QLD 4052 (Doc Data Sheet only)

Intelligence Program

DGSTA Defence Intelligence Organisation
Manager, Information Centre, Defence Intelligence Organisation

Corporate Support Program

Library Manager, DLS-Canberra
MS Sam Doran, Defence Library Service - Sydney West

UNIVERSITIES AND COLLEGES

Australian Defence Force Academy
Library
Head of Aerospace and Mechanical Engineering
Serials Section (M list), Deakin University Library, Geelong, 3217 (
Hargrave Library, Monash University (Doc Data Sheet only)
Librarian, Flinders University

OTHER ORGANISATIONS

NASA (Canberra)
AusInfo
State Library of South Australia

OUTSIDE AUSTRALIA

ABSTRACTING AND INFORMATION ORGANISATIONS

Library, Chemical Abstracts Reference Service
Engineering Societies Library, US
Materials Information, Cambridge Scientific Abstracts, US
Documents Librarian, The Center for Research Libraries, US

INFORMATION EXCHANGE AGREEMENT PARTNERS

Acquisitions Unit, Science Reference and Information Service, UK
Library - Exchange Desk, National Institute of Standards and Technology, US

SPARES (5 copies)

Total number of copies: 54

DEFENCE SCIENCE AND TECHNOLOGY ORGANISATION DOCUMENT CONTROL DATA					
				1. PRIVACY MARKING/CAVEAT (OF DOCUMENT)	
2. TITLE Demand Moderation in Military Communication Networks			3. SECURITY CLASSIFICATION (FOR UNCLASSIFIED REPORTS THAT ARE LIMITED RELEASE USE (L) NEXT TO DOCUMENT CLASSIFICATION) Document (U) Title (U) Abstract (U)		
4. AUTHOR(S) Perry A. Blackmore			5. CORPORATE AUTHOR Electronics and Surveillance Research Laboratory PO Box 1500 Edinburgh SA 5111 Australia		
6a. DSTO NUMBER DSTO-TR-1286		6b. AR NUMBER AR-012-176		6c. TYPE OF REPORT Technical Report	
7. DOCUMENT DATE April 2002					
8. FILE NUMBER E8709/4/15/2		9. TASK NUMBER JTW 99/150		10. TASK SPONSOR DGC4	
				11. NO. OF PAGES 21	
				12. NO. OF REFERENCES 25	
13. URL on the World Wide http://www.dsto.defence.gov.au/corporate/reports/DSTO-TR-1286.pdf				14. RELEASE AUTHORITY Chief, Communications Division	
15. SECONDARY RELEASE STATEMENT OF THIS DOCUMENT <i>Approved for public release</i>					
OVERSEAS ENQUIRIES OUTSIDE STATED LIMITATIONS SHOULD BE REFERRED THROUGH DOCUMENT EXCHANGE, PO BOX 1500, EDINBURGH, SA 5111					
16. DELIBERATE ANNOUNCEMENT No Limitations					
17. CASUAL ANNOUNCEMENT Yes					
18. DEFTEST DESCRIPTORS Communications networks, Military communications, Command control and communications					
19. ABSTRACT Interconnected heterogeneous networks with diverse ownership, carrying a wide range of multimedia traffic with extreme variations in load will characterise military networks of the future. Consequently, the design of a management and control architecture, which fosters efficient and military-tailored resource sharing, is a challenging multifaceted problem. A key feature of such an architecture is to give the network the ability to deliver meaningful signals to users in order for them to modify their behaviour in a way that is beneficial to the network as a whole. For example, during times of network stress, users should be discouraged from excessive network usage. Demand moderation is the term used to encompass the array of mechanisms aimed at achieving this end. Integrated Defence networks of the future should benefit enormously from demand moderation mechanisms.					